



Status of Embedded Linux

March 2024

Tim Bird

Principal Software Engineer, Sony Electronics

Nature of this talk...

- I periodically take a look at the status of embedded Linux
 - Not comprehensive – just stuff I saw
 - I'm sorry if I missed something you're interested in
- Hope to accomplish 2 things:
 - Let people know what's going on
 - Open a discussion on 'what's next' (and what needs more work)

Outline

Linux Kernel
Technology Areas
Industry News
Community
Conclusions

Outline




Linux Kernel

Technology Areas

Industry News

Community

Conclusions



Linux Kernel

- Versions in the last year
 - Pick a few items from each release that are relevant to embedded
 - Lots of things in each release are NOT relevant to embedded (IMHO)
 - Not very good coverage of SoC or driver contributions
- Development Stats
- Some company highlights

Kernel Versions

- Linux v6.3 – 23 Apr 2023 – 63 days
- Linux v6.4 – 25 Jun 2023 – 63 days
- Linux v6.5 – 27 Aug 2023 – 63 days
- Linux v6.6 – 29 Oct 2023 – 63 days
- Linux v6.7 – 7 Jan 2024 – 70 days
- Linux v6.8-rc7
 - Expect 6.8 on 17 March 2024 (or maybe earlier)

Linux v6.3 (April 2023)

- Lots of old unused Arm board files were removed
- Kernel can be configured with a built-in Dhrystone test
- Default “make V=0” option has been removed
- Minor change to Developer’s Certificate of Origin clarifies that nicknames can be used for contributions
 - As long as your identity is known
 - See <https://git.kernel.org/pub/scm/linux/kernel/git/torvalds/linux.git/commit/?id=d4563201f33a>

Linux v6.4 (June 2023)

- The SLOB memory allocator was removed
- Some nice documentation added for building the kernel
 - <https://docs.kernel.org/admin-guide/quickly-build-trimmed-linux.html>
- MODULE_LICENSE() declarations were removed, for code that cannot be built as a module
 - See <https://lwn.net/Articles/927569/>
- User trace events (API fixes) was merged
 - See <https://lwn.net/Articles/927595/>

Linux v6.5 (Aug 2023)

- Parallelize bringing up CPUs, reducing time to get all processors online
 - But only on X86, for now
- There was a big reorganization of 32-bit devicetree file directory structure
 - To match arm64, mips and riscv
 - See commit 6c1561fb
- Miscellaneous ftrace and perf updates
 - See <https://lwn.net/Articles/937006/>

Linux v6.6 (Oct 2023)

- Can now build a kernel without buffer-head support
 - Some buffer head functionality can be replaced by iomaps (and folios)
 - Part of an effort to eliminate the buffer head subsystem
 - See <https://lwn.net/Articles/931809/>
 - This might lead to older filesystems (that depend on the buffer cache) being deprecated
- New 'eventfs' subsystem
 - Improves memory efficiency of the tracing subsystem
 - Only allocates inodes and dentries for tracepoints when they are needed
- Shadow Stacks (on Intel, for now)
 - See next page

Shadow Stacks

- Provides an alternate stack for return addresses (only)
 - Return address is put on regular stack AND on shadow stack
 - Often managed by hardware, for speed (and security)
 - On return, shadow stack is used to detect if the return address on the regular stack has been altered
 - VERY useful for detecting return-oriented programming (ROP) attacks, which are a major source of security threats
- x86 support landed in 6.6
- Patches for 64-bit ARM and RISC-V are in development
- See <https://lwn.net/Articles/940403/>
- Work on shadow stack support in MUSL C library has started

Linux v6.7

- Can enable or disable 32-bit emulation on x86-64 kernels on the kernel command line
 - Can reduce attack surface
- ia64 (Itanium) architecture removed
- More efficient kernel samepage merging
 - See <https://docs.kernel.org/admin-guide/mm/ksm.html>
- Printk now supports per-console locking
 - Allows high-priority message output to take over console from lower-priority message
 - <https://lwn.net/Articles/909980/>

Linux v6.8-rc7 (6.8 Expected March 2024)

- 'deadline servers' added
 - Prevents realtime tasks from completely starving normal tasks
- The first real RUST driver was added to the kernel
 - A reference driver for Asix PHYs
- Networking core data structures were refactored for better cache efficiency
 - Improving TCP performance in some situations up to 40%
- The SLAB memory allocator was removed (!!)
 - SLUB is the last one remaining in the kernel
- perf data-type profiling was added

deadline servers (for less deadly RT)

- Provides a way to prevent realtime tasks from starving normal tasks
 - Allows an admin to fix a "runaway realtime task" problem
 - Replaces realtime throttling feature
- Added in Linux v6.8
- How it works:
 - Deadline scheduling priority is higher than RT priority
 - A deadline server is set up to be started when possible normal task starvation is detected
 - The deadline server runs normal tasks, within it's guaranteed scheduling quanta (e.g. 5%)
- See <https://lwn.net/Articles/934415/>

perf data-type profiling

- Adds a new sub-command to perf
 - `perf annotate --data-type`
 - `perf annotate --data-type={field-name}`
- Can get sample data per-field in a data structure
- Shows offset of structure fields and frequency of access
- Is very useful to see if shuffling the fields in a data structure might (or did) improve cache performance
 - Can check if a frequently accessed field is in a different cache line
 - Or if all frequently accessed fields are in the same cache line
- See <https://lwn.net/Articles/955709/>

Linux community issues

- Recent discussions about the threat of maintainer burnout
 - This has been a long-term problem
 - Workload is too high, and turnover is very slow
 - Not enough assistant or replacement maintainers
 - Mental health session was provided at LPC for dealing with maintainer stress
 - Also a session with kernel leaders discussing maintainer burnout
 - See <https://lwn.net/Articles/952034/>
- Lots of kernel mailing lists moved from lists.linux-foundation.org domain to lists.linux.dev
 - Should adjust your address book for kernel mailing lists

Linux 6.6 developer stats

- 14,069 change sets, by 1978 developers (249 new devs.)
- Most active 6.6 developers, by changesets:

Person	Org.	Changesets	Percent	Subsystem Area
Krzysztof Kozlowski	Linaro	311	2.2%	device tree updates
Yangtao Li	Vivo	296	2.1%	converted lots of drivers to use devm_platform_iomap_resource()
Konrad Dybcio	Linaro	188	1.3%	device tree files and Qualcomm drivers
Rob Herring	Arm	184	1.3%	device-tree related code
Thomas Zimmermann	Suse	158	1.1%	graphics and frame-buffer drives

Linux 6.6 developer stats

- Most active 6.6 developers
- By lines of code changed:

Person	Org,	Lines changed	Percent	Subsystem area
Jiri Pirko	? (unaffiliated?)	17444	3.0%	Refactor of networking subsystem
Konrad Dybcio	Linaro	16477	2.8%	device-tree and Qualcomm drivers
Ian Rogers	Google	8991	1.5%	Lots of perf cleanups
Dmitry Baryshkov	Linaro	8729	1.5%	Qualcomm driver work
Charles Keepax	Cirrus Logic	7834	1.3%	Intel sound driver work

Kernel commit log entries

- Number of commit log entries (including merges), per kernel version

Company	git log count	developer count
6.3	15637	2052
6.4	16011	2076
6.5	14674	2008
6.6	15289	2045
6.7	18404	2059
6.8-rc7	15398*	1980*

- Extracted using 'git log v6.yy..v6.zz --oneline | wc -l' and 'author-stats v6.yy..v6.zz | wc -l'

Contributions by embedded Linux companies (to kernel)

Company	Commits since April 2023	Top contributor	Work area(s) (of top contributor)
Baylibre	207	David Lechner	AD2S1210 driver, spi stuff
Bootlin	373	Miguel Raynal	mtd rawnand, nvmmem, mac802154 driver
Collabora	774	AngeloGioacchino Del Regno	mediatek processor support, panfrost GPU
Ideas On Board	358	Laurent Pinchart	12c, camera, media drivers
Igalia	110	Maíra Canal	drm GPU driver
Linaro	4681	Krzysztof Kozlowski	sound, device tree, Samsung clocks
Linutronix	458	Thomas Gleixner	timers, printk, x86, preempt_rt
Pengutronix	2519	Uwe Kleine-König	driver cleanups (remove callback returning void)
Toradex	112	Francesco Dolcini	drm bridge fixes
Wind River	67	Ovidiu Panait	sahara crypto driver

Long Term Supported Kernels

- 6.6 is the latest long-term support (LTS) kernel
 - Current LTS kernels are: 4.14, 4.19, 5.4, 5.10, 5.15, 6.1, and now 6.6
- Plan is to pick one per year, and only maintain for 2 years
 - This is a reduction from the previous plan to maintain for 6 years
 - As support expires for older LTS kernels, they will not be replaced as soon
 - Will reduce the number of supported LTS kernels from 6 to 2
- If you need longer support, get kernels from CIP or Ubuntu LTS
 - Both of these plan to support kernels for 10 years
- See <https://www.zdnet.com/article/linux-kernel-6-6-is-the-next-long-term-support-release/>

Kernel releases (as of March 2024)

The Linux Kernel Archives



- About
- Contact us
- FAQ
- Releases
- Signatures
- Site news

Protocol	Location
HTTP	https://www.kernel.org/pub/
GIT	https://git.kernel.org/
RSYNC	rsync://rsync.kernel.org/pub/

Latest Release

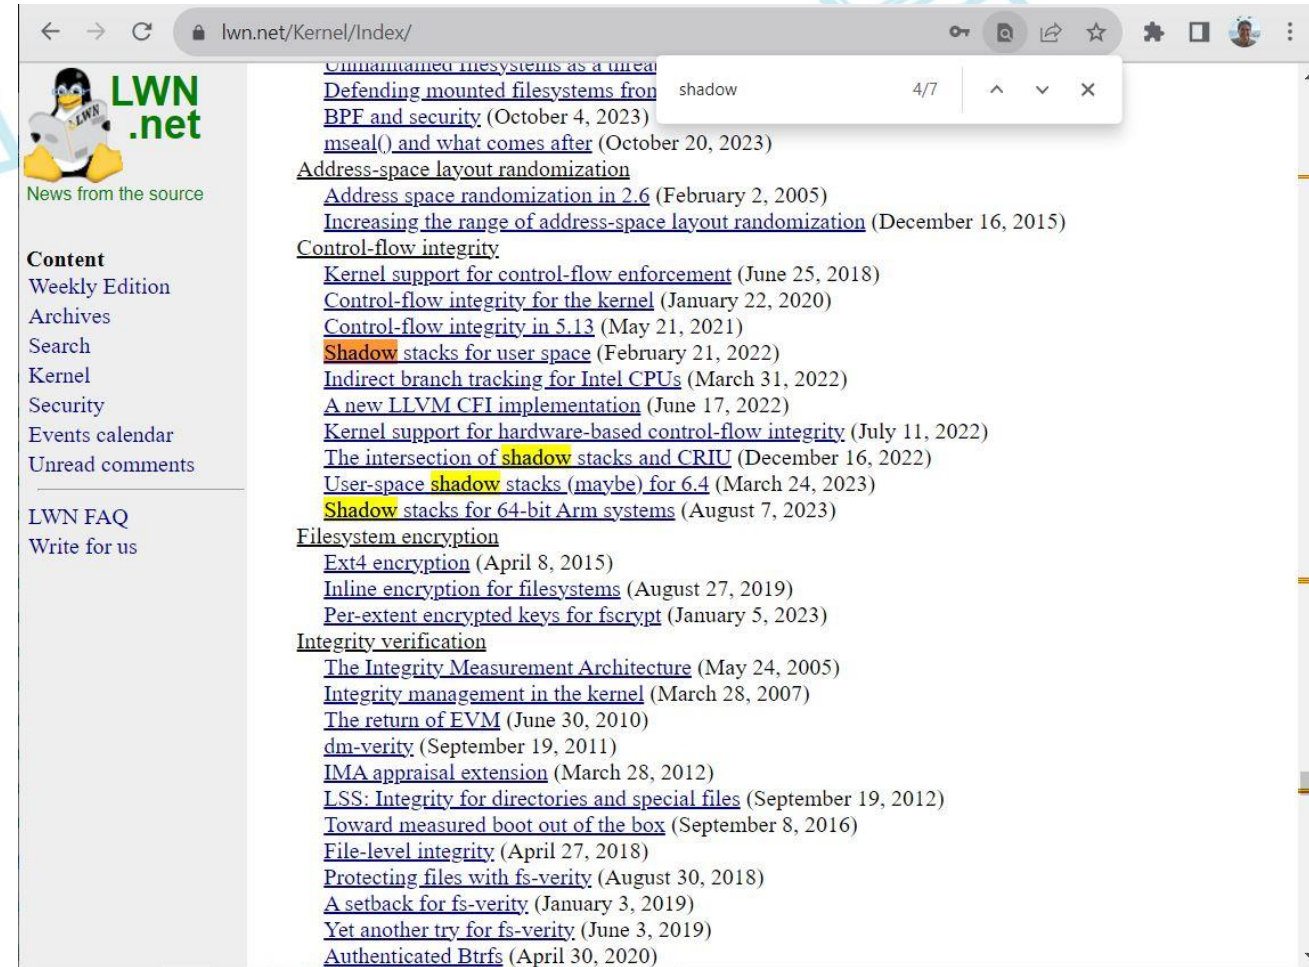
6.7.7

Download icon

mainline:	6.8-rc6	2024-02-25	[tarball]	[patch]	[inc. patch]	[view diff]	[browse]	
stable:	6.7.7	2024-03-01	[tarball]	[pgp]	[patch]	[inc. patch]	[view diff]	[browse] [changelog]
longterm:	6.6.19	2024-03-01	[tarball]	[pgp]	[patch]	[inc. patch]	[view diff]	[browse] [changelog]
longterm:	6.1.80	2024-03-01	[tarball]	[pgp]	[patch]	[inc. patch]	[view diff]	[browse] [changelog]
longterm:	5.15.150	2024-03-01	[tarball]	[pgp]	[patch]	[inc. patch]	[view diff]	[browse] [changelog]
longterm:	5.10.211	2024-03-01	[tarball]	[pgp]	[patch]	[inc. patch]	[view diff]	[browse] [changelog]
longterm:	5.4.270	2024-03-01	[tarball]	[pgp]	[patch]	[inc. patch]	[view diff]	[browse] [changelog]
longterm:	4.19.308	2024-03-01	[tarball]	[pgp]	[patch]	[inc. patch]	[view diff]	[browse] [changelog]
linux-next:	next-20240301	2024-03-01						[browse]




LWN.net Resource for Kernel Information

- LWN.net Kernel Index
 - Has many years of articles
 - Sorted by kernel subsystem and topic
- <https://lwn.net/Kernel/Index/>



Outline

Linux Kernel
Technology Areas
Industry News
Community
Conclusions



Technology Areas



- Architectures
- Bootloaders
- Boot Time
- Core Kernel
- Filesystems
- Languages
- Networking
- Security
- Testing
- Toolchains
- Tracing
- System Size
- Build Tools and Distros

Architectures

- LoongArch support recently added (v5.19 and 6.x kernels)
- RISC-V support is becoming very mature
- Support for many old ARM boards was recently removed
- Removal of old Architectures
 - Super-H
 - Itanium

Bootloaders

- U-Boot
 - Now supports loading images over HTTP
 - Previously only supported the UDP protocol
 - Could only use NFS or TFTP as servers
 - Now can download kernel and other images (dtb, initrd, etc.) from a web server (e.g. using wget)
 - See <https://www.linaro.org/blog/http-now-supported-in-u-boot/>
 - Includes new "binman" tool for packaging and managing firmware images
 - Improved support for tracing, firmware handoff, verified boot, LTO
 - Better docs and CI system
 - See "Status update for U-boot" (ELC 2023 session by Simon Glass)
 - <https://lwn.net/Articles/938769/>

snagboot

- Snagboot is a set of tools that can help boot and install images on boards that fail to boot
- Produced by Bootlin
- Consists of
 - snagrecover – to initialize memory and run your bootloader
 - snagflash – to flash a working system image using either DFU, USB Mass Storage, or fastboot
- Works on several different boards, and replaces proprietary tools
- See <https://bootlin.com/blog/releasing-snapboot-a-cross-vendor-recovery-tool-for-embedded-platforms/>

Boot Time reduction


- Resurrection of Ureadahead (as an upstream project)
 - Ureadahead allows to pull pages into memory so they are already in the page cache when needed by a process
 - Has been used for years by Chromebook team (since 2009), but upstream went inactive
 - Want to revive public, upstream project
 - Presentation at ELC 2023:
https://elinux.org/images/5/5b/Ureadahead_resurrection.pdf
 - by Steve Rostedt
 - Video - <https://youtu.be/HwdWKMxM83E>

TI Boot time ideas

- Texas Instruments (TI) presentation on automotive boot time
- Ideas:
 - Custom bootloader to avoid duplicate device initialization
 - But must support safety requirements
 - Maybe parallelize parts of u-boot
 - Specialized MCUs are used to start hardware before Linux boots
 - But then a handoff is needed from MCU firmware to Linux
 - Would be nice to standardize control of such MCUs from Linux
 - Open ethernet MAC in driver probe
- Presentation: "Resolve and standardize early access to hardware for automotive industry with Linux"
 - by Khasim Syed Mohammed at LPC 2023
 - See <https://lpc.events/event/17/contributions/1499/>

Core Kernel

- Improvements to kernel same-page merging
 - Smart scanning, and auto-tuning (coming), for better efficiency
 - See <https://lwn.net/Articles/953141/>



Graphics – GPU work





- Open Source Vulkan driver for NVIDIA hardware
 - Read for prime time
 - Announced by Collabora
 - See <https://www.collabora.com/news-and-blog/news-and-events/nvk-is-now-ready-for-prime-time.html>

Filesystems and I/O

- Some work on MTD spi-nor
 - Enhanced locking to support reads while writing
- EROFS enhancements
 - Caching and speed improvements
- Block IO speedup (coming in v6.9?)
 - Cache time to avoid lookup overhead during issue-side of IO ops
 - Yields about 6% performance improvement
 - See <https://www.phoronix.com/news/Linux-Caching-Time-Block-IO>

Languages

- Python
- Rust



Languages - Python

- Python 3.11 (April 2023) had a lot of improvements
 - Better tracebacks - Fine-grained error locations in tracebacks
 - Faster
 - Claims a 1.22x speedup on the standard benchmark suite
 - Up to 10-60% faster than Python 3.10
 - Exception groups and except*
- Python 3.12 released in October, 2023
 - <https://docs.python.org/3.12/whatsnew/3.12.html>
 - Python 3.12 improvements
 - Typing improvements
 - os and pathlib performance improvements
- Python 3.13 will have a JIT (Just-in-time compiler)
 - See <https://tonybaloney.github.io/posts/python-gets-a-jit.html>

Languages - Rust

- Rust support continues to go into the mainline kernel
 - Finally have an example of a "real" driver in 6.8
 - Still only an example
 - Also have a (RFC for a) Rust implementation of Android's Binder (IPC mechanism)
 - See <https://lwn.net/Articles/953116/>
 - Still a big debate about whether this is positive or negative for future maintenance
 - See <https://lwn.net/Articles/952029/>
- Rust 1.76 released on Feb 8, 2024

Networking

- Many industries are moving from proprietary networking to ethernet
 - Many are using single-wire (SPE = Single Pair Ethernet)
 - Reduces weight and amount of copper needed
 - 10/100/1000/.../BaseT1/L/S (vs 10/100/100/BaseT)
 - *S is used in automotive, and supports multi-drop, single-pair, cabling
 - Many uses of PoE (Power over Ethernet) and PoDL (Power over Data Line)
 - Eliminates separate power cable
- Presentation: "Braiding wires into the Linux network stack"
 - by Oleksij Rempel, at ELC 2023
 - Slides: [https://elinux.org/images/7/7e/EOSS - 2022 - embedded networking - 2023.06.20.pdf](https://elinux.org/images/7/7e/EOSS_-_2022_-_embedded_networking_-_2023.06.20.pdf)

Real-Time

- PREEMPT_RT status
 - Patches have been going in continuously – through 6.7
 - Although some stalled (like the printk refactor)
- Xenomai
 - The other Linux realtime approach (dual-kernel)

PREEMPT_RT - What's left

- What's left in PREEMPT_RT patches out of mainline:
 - last year (2023) (patches-6.7-rc3-rt3):
 - About 2800 lines of code, affecting 79 files (in 91 patches)
 - Some changes to the printk, the core scheduler, some locking and timer tweaks, DRM, and a few other places.
 - this year (2024) (patches-6.8-rc7-rt6):
 - about 2800 lines of code, affecting 87 files (in 118 patches)
 - 52 printk patches, 18 tty/serial patches, 12 drm, 5 scheduler
 - A few patches each for ARM, PowerPC, Risc-V
 - Thomas said at Plumbers that printk changes are the blocker to the patch that allows enabling PREEMPT_RT in the mainline kernel
- See <https://mirrors.edge.kernel.org/pub/Linux/kernel/projects/rt/6.7/>

PREEMPT_RT – how to find detailed status

- Realtime Linux workshop at Embedded Open Source Summit 2023
 - Good talks about current issues:
 - RTLA timelat, cyclictst latency, new tracer to monitor RT task behavior, etc.
 - See YouTube playlist for event: <https://youtu.be/NWVWXtfOzXM>
 - Detailed status of PREEMPT_RT work: <https://lwn.net/Articles/938236/>
- "Evaluation of PREEMPT_RT in Virtualized Environments"
 - Presentation by Jan Altenberg at ELC 2023
 - See https://elinux.org/images/0/05/Preempt_rt_virtualization.pdf
- Detailed status from plumbers by Thomas Gleixner
 - "RT preemption end game"
 - See <https://lwn.net/Articles/951337/>
 - TL;DR – Need to get refactored printk work upstream

Xenomai

- Xenomai = dual-kernel hard realtime scheduling approach
 - Realtime runs outside of Linux with a separate scheduler
- Xenomai has 2 main versions:
 - Version 3.2
 - Out-of-source-tree: Cobalt scheduler core + libraries
 - Supports POSIX natively, emulates VxWorks, PSOS and “Native” APIs
 - Drivers are forks/rewrites against the Real-Time Driver Model (RTDM)
 - Version 4.0
 - EVL (EVEN Less) Core - comes as kernel patch
 - Supports a new ABI - No POSIX, no legacy RTOS emulation
 - Drivers are kernel drivers with patched-in “Out-ofBand” support
- Presentation: "Building Hard real-time systems with Xenomai"
 - by Jan Kiszka, at OSSJ 2023
 - See <https://sched.co/1TyqT>
 - Discusses "When to use PREEMPT_RT and when to use Xenomai"?

Security

- Shadow Stacks
 - New in v6.6 (already discussed)
- Work to remove ELF notes from sysfs
 - Some ELF notes in sysfs included relocated kernel addresses
 - Which defeated KASLR (Kernel Address randomization)
 - See <https://lwn.net/Articles/962782/>
- BPF raises some interesting security issues (unsurprisingly)
 - Authoritative LSM hooks (ones that can override other security mechanisms) was proposed for BPF, but rejected
 - There has been talk of allowing unprivileged users to run BPF modules
 - Not sure if anything will happen there
 - See <https://lwn.net/Articles/929746/>

System Size

- SLOB memory allocator removed from kernel (v6.4)
 - Not enough people using it (or reporting that they use it)
 - Was deprecated in v6.2
 - IMHO, 2 releases (5 months) isn't enough time for people to notice the deprecation and object to it
- SLAB memory allocator removed from kernel (v6.8..)
- Only SLUB remains
 - Use CONFIG_SLUB_TINY for small systems

Testing

- Suites of tests
- New proposal for test declaration per kernel sub-system
- Proposal for gitlab-ci testing pipelines in upstream

Test Suites

- LTP
 - Linux Test Project
- kselftest
 - Lots of tests and test improvements for bpf, mm, resctrl, hid, xsk, networking, x86 features, powerpc, KVM, landlock and more
- Kunit
 - kunit tool improvements
 - Proposal to support device testing (see subsequent page)

Test Suites - LTP

- LTP = Linux Test Project
- Big project to test kernel from user-space
 - Maintained by SUSE
 - Consists of POSIX, Realtime, system-call tests, and others
- Latest release = LTP 20240129 release
 - New tests and fixes to tests
 - New "runltp-ng" renamed to "kirk"
 - New test execution manager (still experimental)
 - Experimental LTX program (from LTP 20230929 release)
 - Tiny binary that executes tests on a separate machine
 - See <https://github.com/linux-test-project/ltp/releases>

Test Suites - kselftest

- kselftest/kunit proposed changes to KTAP
 - KTAP = Kernel Test Anything Protocol
 - Is the standardized output for kernel tests
- Recent proposal to add meta-data to tests and KTAPv2
 - Kunit Attributes: set of well-known tags
 - speed (fast or slow) – indicating speed category of test
 - is_init – indicates scope of execution
 - test_params – indicates parameters used for test
 - Can filter tests at runtime or results at analysis-time
- See "Storing and Outputting Test Information: Kunit Attributes and KTAPv2" (Plumbers 2023 presentation)
 - See <https://lpc.events/event/17/contributions/1527/>

Test Suites - Kunit

- Recent proposal is for testing of devices
- Check results of probing devices listed in device-tree
- Check results of instantiating devices on discoverable busses
 - Tests attributes in sysfs for devices, using a per-board definition file
 - Is first hardware-specific testing proposal I've seen upstream
 - Requires a set of board/device mappings, that upstream doesn't want to maintain
 - Probably develop an out-of-tree repository for this

Test declaration for kernel sub-systems

- Proposal for new system to specify a test for a sub-system
 - Add "V: <test-recommendation>" field to MAINTAINERS file
 - Can specify a test, or point to documentation about test to run
 - Suggested by Nikolai Kondrashov (RedHat testing group)
 - Still under development
- Eventually, may require patches to include attribute:
 - Tested-with: ...
 - Could include just test name, or reference to test results
- Goal is to allow:
 - Contributors to more easily test their patches
 - Maintainers can verify that patch doesn't break things

gitlab-ci testing pipelines

- Collabora proposes some gitlab-ci testing materials
 - Submitted patches to Linux kernel In January and February 2024
- Interesting set of gitlab pipeline files and shell scripts
- Currently supports static checks (smatch and checkpatch)
- If you are using gitlab, this could be quite interesting for your QA department
- See <https://www.collabora.com/news-and-blog/news-and-events/patch-submitted-to-introduce-gitlab-ci-pipeline-for-kernel-testing.html>

Toolchains - GCC

- GCC 13.2 released July 27, 2023
 - See <https://gcc.gnu.org/gcc-13/changes.html>
 - LTO (link-time optimizations) has been improved
 - Can now emit diagnostics in the SARIF format and gcc's own JSON-based format
 - SARIF = Static Analysis Results Interchange Format
 - Support for C23 features
 - LoongArch support improvements
 - Improvements in the static analyzer

Toolchains - LLVM

- LLVM 17.0.6 released November 28, 2023
 - See <https://releases.lvm.org/17.0.1/docs/ReleaseNotes.html>
 - CLANG can output multiple lines of source on errors
 - Diagnostics for pre-processor errors and attribute errors
 - Some warnings changed to match GCC's behaviour
 - Many other diagnostic message improvements during compilation
 - See <https://blog.lvm.org/posts/2023-09-19-diagnostic-improvements-in-clang-17/>
- People are using LLVM for whole distributions, not just the kernel

Toolchain resources

- Presentation: "Toolchain Options in 2023: What's New in Compilers and Libcs"
 - at ELC 2023, by Bernard Rosenkranzer, BayLibre
- Still can use traditional toolchain elements, but there are interesting alternatives:
 - binutils: ld, lld (LLVM), mold linker
 - Mold linker is intended to be much faster (and it is)
 - Support LTO for both gcc and LLVM
 - Compilers: gcc, clang (LLVM), TinyCC
 - Similar performance, some BSPs only on gcc or clang
- See [https://elinux.org/images/8/84/EOSS23 - Toolchain Options in 2023 - What%27s New in Compilers and Libcs.pdf](https://elinux.org/images/8/84/EOSS23_-_Toolchain_Options_in_2023_-_What%27s_New_in_Compilers_and_Libcs.pdf) and <https://youtu.be/Vgm3GJ2ItDA>

Tracing

- Something is always going on with perf
 - New features seem to be added every release
 - Example: perf data-type tracing (covered previously)
- eBPF
 - Introduction, impact and status (new LF report):
 - <https://www.linuxfoundation.org/research/state-of-ebpf>
- eBPF is now used for tracing and monitoring (including itself)
 - Netflix releases bpftop for monitoring bpf programs
 - See <https://thenewstack.io/netflix-releases-bpftop-an-ebpf-based-application-monitor/>

Build Tools and Distros




- Yocto Project
 - Latest version = 4.3 (Nanbield), released November, 2023
 - Kernel 6.5 and 300+ recipe upgrades, LLVM 17 support
 - runqemu enhancements
 - Lots of security fixes
 - Now requires Python 3.8
 - Bunch of other changes
 - See release notes:
 - <https://downloads.yoctoproject.org/releases/yocto/yocto-4.3/RELEASENOTES>

Kernel community – CVE handling

- Linux kernel community becomes a CNA
 - CVE assignment is fraught with issues
 - Sometimes the project members disagrees with the impact and severity assessments of the CVE submitter
 - CAN (Certified Number Authorities) have input into this process
 - See <http://www.kroah.com/log/blog/2024/02/13/linux-is-a-cna/>
 - Currently Lee Jones, Sasha Levine and Greg Kroah-Hartman
 - Documentation about the process CVE/CAN process are in this patch:
 - <https://lore.kernel.org/lkml/2024021430-blanching-spotter-c7c8@gregkh/>
- Note: glibc project also became a CNA

Outline

Linux Kernel
Technology Areas
Industry News
Community
Conclusions



Industry News

- Legal issues
 - Update on SFC vs Vizio
- OpenWRT project designing its own hardware
- HDMI Forum preventing creation of open source HDMI driver
- Interesting uses of Linux
 - OSS in space!



Update on SFC vs. Vizio

- Lawsuit is still ongoing
- Is important, because SFC is asserting a novel interpretation of GPL rights
 - 3rd party beneficiary rights
 - which would allow any 3rd party to enforce the license
 - GPL is a contract, not a license
- Latest update:
 - Judge denied Vizio motion for summary judgement (Jan 3, 2024)
 - Appears to interpret GPL as a contract (as SFC asserts)
 - no "preemption" by copyright law (no removal to federal court)
 - See <https://sfconservancy.org/news/2024/jan/03/vizio-sj-rejected/>
 - Trial likely in June, 2024
- Random note: Walmart is trying to acquire Vizio

OpenWRT project designs its own hardware



- OpenWRT originally based on Linksys router software
 - Project is now 20 years old
- Project wants standard reference hardware
 - Making "OpenWRT One" board, with help from banana-pi org.
- Wants to make a fully open system
 - Users not stuck with proprietary firmware, anywhere in system
- See <https://arstechnica.com/gadgets/2024/01/openwrt-now-20-years-old-is-crafting-its-own-future-proof-reference-hardware/>

HDMI 2.1 can't have an open source driver

- HDMI Forum told AMD that it can't make an open source driver for HDMI 2.1
 - Means that Linux users can't use highest resolution and frame rates
 - Is a big problem for GPU drivers (which often include audio driver technology for HDMI support)
- Public access to spec has been a problem since at least 2021
 - See <https://www.phoronix.com/news/HDMI-Closed-Spec-Hurts-Open>
- See <https://arstechnica.com/gadgets/2024/02/hdmi-forum-to-amd-no-you-cant-make-an-open-source-hdmi-2-1-driver/>

Interesting embedded Linux uses

- Satellites
 - CubeSats
 - Starlink satellite constellation
- Mars Ingenuity helicopter
- OSS on the moon!




Linux in Satellites

- Linux is currently being used in many satellites
- CubeSats
 - Linux has been used in satellites since 2003
 - Lots of experiments with Linux and COTS in cubesats since 2014
 - One example: NASA PhoneSat (using an Android phone as the flight computer for a cubesat)
 - One estimate is that about 50% of cubesats run Linux (in some part of the flight stack)
- Major constellations (StarLink and Planet) use Linux
 - (see next slide)

Starlink Satellite constellation

- Each new Starlink satellite uses over 80 processors, each one running Linux
 - Uses clusters for fault tolerance
- There are now (as of March 2024) over 5500 Starlink satellites currently in orbit
 - Over 400,000 nodes of Linux in orbit!
 - <https://satellitemap.space/> is quite interesting!!




Starlink fleet upgrade information

- Very interesting presentation at SmallSat 2023 conference
- Uses A/B boot system
 - Can always revert to a known working image
- Single image uploaded to a master processor in each satellite
 - Update image contains sub-images for other processors and components
- Automated, but controlled
 - Ground station controls configuration
 - Satellite automatically update via https
 - Can do an update to a set of "canary" satellites, to find any problems
- See <https://digitalcommons.usu.edu/smallsat/2023/all2023/72/>
 - Slides:
<https://digitalcommons.usu.edu/cgi/viewcontent.cgi?filename=0&article=5598&context=smallsat&type=additional>

Falcon 9 and Starship Rockets

- Falcon 9 uses Linux for avionics
- Starship (prototype next-gen reusable rocket) is presumed to use the same avionics stack




Mars Helicopter - Ingenuity




Mars Helicopter

- Mars Ingenuity Helicopter landed in February, 2021 on Mars
- Helicopter blades were damaged on flight 72, on January 18, 2024
 - Helicopter can no longer fly
- NASA announced the end of the mission






OSS on the Moon!

- Sora-Q transformable Lunar Rover (Also known as LEV-2)
 - Landed on the moon on the SLIM mission, by JAXA, on Jan 19, 2024
 - Developed by JAXA, Tomy, Sony, and Doshisha University
- Took this awesome picture of the SLIM lander
- OK – it's not Linux, but NuttX (but that's still open source)






Outline

Linux Kernel
Technology Areas
Industry News
Community
Conclusions



Community

- Conferences
- Elinux wiki
- Trade Associations and Projects



Conferences

- Embedded Linux Conference
 - Is now back to twice a year
 - In the spring (April/May) in North America and fall in Europe
- Last one = June 28-30, 2023 in Prague, Czech Republic
- Next one = April 16-18, 2024 in Seattle, USA
- Following one = Sept. 16-17, 2024 in Vienna, Austria
 - Will be a smaller event this year
 - 2 days, 2 tracks
- Embedded micro-conference at Linux Plumbers
 - Also in Vienna, Sept. 18-20, 2024

ELC 2023 Information

- June 28-30 in Prague, Czech Republic
- Part of Embedded Open Source Summit
- Focus areas:
 - Linux in Space and aerospace
 - micro-track on AOSP (Android)
 - Realtime
 - boot time, video & camera sub-system, new boards
 - And lots more (over 60 talks and sessions)
- See https://elinux.org/ELC_Europe_2023_Presentations

Conferences (recent)

- Embedded Open Source Summit 2023
 - June 27-30, in Prague
 - Embedded Linux Leadership Meeting (June 30)
- Japan Technical Jamboree (episode 2)
 - September 29, in Tokyo
- Embedded Recipes
 - September 28-29, in Paris
- Linux Plumbers
 - November 13-15, 2023 in Richmond, Virginia, USA
 - Had an Embedded BOF

More Conferences

- FOSDEM
 - February 3-4 2024 in Brussels, Belgium
 - Embedded track:
<https://fosdem.org/2024/schedule/track/embedded-mobile-and-automotive/>
- Linaro Connect
 - May 14-17, 2024 in Madrid, Spain
- Linux Plumbers
 - September 18-20, 2024
 - Plan to have an embedded Linux micro-conference

Embedded Linux Leadership Summit 2023

- Meeting of Embedded Linux ecosystem stakeholders
- Invitation-only event, with:
 - Product companies (Sony, Google, Amazon)
 - Processor vendors (ARM)
 - Trade associations (LF, Linaro)
 - Project leaders (AGL, CIP, KernelCI, Zephyr, RT Linux, Yocto)
 - Consulting companies (Baylibre, Collabora, Bootlin, Pengutronix)
- Discuss status of ecosystem, and technology areas that need work

Issues From Embedded Linux Leadership Summit 2023

- Security is an ongoing concern
- OTA updates
- Long-term support
- Heterogeneous core support
- Shared testing
- Safety critical / certification
- Standardization of HW components
- AOSP competition
- Lack of investment in infrastructure/ tools/ upstreaming
 - (not enough contributions)
- RT marketing
- Upstream AI/ML accelerators
 - Support for AI processing units (APU)

Plumbers Embedded Linux BOF

- Session at Linux Plumbers Conference to discuss issues with embedded Linux
 - Slides: <https://lpc.events/event/17/contributions/1549/attachments/1344/2692/Embedded-Linux-BOF-2023-11-LPC.pdf>
 - Video: <https://youtu.be/1sIW64Qip-I>
- Summary of major topics discuss in next few slides

Technology Areas



Technology Areas

- Architectures
- Bootloaders
- Boot time
- Filesystems
- Networking
- Security
- Testing
- Tools
- Toolchains
- Tracing
- System Size
- Build Tools and Distros

- Attendees said that "Video" and "Power Management" should be included in list of technology areas.

Survey of "most interested" areas

Topic Area	Total Votes
Architecture	7
Bootloaders	4
Boot time	4
Filesystems	0
Networking	2
Security	4
Testing	6
Tracing	1
System Size	3
Build Tools and Distros	5*

Bootloaders and Boot Time

- Resurrection of Ureadahead
- TI LPC presentation on automotive boot time
 - Already covered (in this presentation) previously

Boot time discussion topics

- Kernel Execute-in-Place (XIP)
 - Sony used XIP a long time ago, but it seems dead upstream now
 - Linaro did some work to support XIP, but nobody seems to be using it
- Asynchronous probing
 - Seems underutilized upstream
 - Needs a lot of research to validate that drivers come up properly
- Modularity and deferred driver loading

Boot time discussion topics

- Coming out of low power state (suspend/resume)
 - Use Suspend/Resume (or unhibernate) instead of cold boot
- Android has a slow cold boot, and could use resume from flash (rather than cold boot from flash)
- Device links in device tree helps reduce boot time by lowering the number of deferred probes
- TI divided a network driver into early and late initialization
- Loading and authentication of firmware takes a lot of time
 - Would be nice to offload this (separate it from Linux startup)

Why is embedded so hard in open source?

- OSS means finding generalizations that can apply to as many users/products as possible
- Embedded means specializing (customizing) the software for your product
- Lots of development areas in embedded, like system size, boot time, and power efficiency require system-wide specializations
 - There's no single area you can attack to meet your requirements
 - Requires "subtractive engineering"

File Systems

- Is anyone still using MTD devices (raw flash)?
 - Yes, almost exclusively with UBIFS
- Everyone else seems to be using ext4 on MMC devices
- Not a lot of work on embedded-specific filesystems
 - Although f2fs, cramfs and other still being worked on

System Size

- Tim upset that SLOB allocator was deprecated so quickly
 - And patches to remove SLAB are on the mailing list
- Arnd asks if 32-bit is being less used
 - But everyone in the room said that they have used 32-bit in the last year
 - For new processors, there's a huge trend to move away from 32-bit
- Linux fits just fine in smallest purchasable DRAM
 - Other Oses (such as Zephyr) should be used for sub-8M systems
 - Linux has lost markets where only SRAM is used (sub-2M)
- RISC-V supports 32-bit (but with smallest memory of 64M)

System Size

- How many 32-bit users are using more than 4GB of memory?
 - Some people are using it, but upstream wants to remove high memory for 32-bit systems
- Nolibc is an interesting alternative to a C library
 - Can it be used for small-system development

Build tools and Distros




- Yocto Project
 - Had a release since
 - Would like to see more frequent LTS releases for kernel in YP
 - YP can only afford to do an LTS release every 2 years
 - Not enough resources to do it more often
- Not enough resources for Yocto Project (and other embedded Linux projects) in general
 - Companies using Yocto should invest more

Issues from Embedded Linux Leadership Summit 2023

- Ecosystem around cameras and AI needs a lot of work
- Android has a standardized API for lots of things
 - Google provides a lot of resources for their ecosystem

Community issues

- Mailing lists
- elinux Wiki
- Conferences



Community issues – mailing lists






Embedded Linux email lists

- For kernel code:
 - linux-embedded@vger.kernel.org
 - Status: not used very much, never suggested by `get_maintainers.pl`
 - The last substantive discussion was about the MAINTAINERS file, in 2021
- For strategy and community-building discussions:
 - elinux-discuss@lists.elinux.org
 - Status: brand new

Community issues – elinux wiki

- Discussed later...



Community issues – conferences



Events - ELC

- Embedded Linux Conference
 - Is now only once a year
 - Alternates between Europe and North America
 - So, effectively only once per region every two years
 - May decide to fill in with regional events or Plumbers
 - Sometimes in Embedded Open Source Summit, and sometimes in Open Source Summit (North America)
 - Goal = Consistently in the spring (April/May)
 - next one: 2024 = April 15-19 in Seattle, USA

Open Discussion – what's needed?

- Need a group of maintainers to help people
- SoC vendors are doing more upstreaming
 - Maintainers need help to keep up
- Would be good to have documented design



<< End of Plumbers Embedded Linux BOF >>

Elinux wiki

- Linux Foundation hired Bill Traynor as a fulltime employee
 - Will do elinux administration as part of his job
- Site is still used for:
 - Materials for embedded Linux development boards
 - Some academics use it for coursework
 - Event materials: slides and links to videos for ELC
- Some areas of the site are out-of-date
- The site is underutilized for sharing information
- Looking for volunteers to help with the site
- Have set up an LFX Crowdfunding site
 - <https://crowdfunding.lfx.linuxfoundation.org/initiative/5fa1a40a-d4f0-4576-81c6-57319a591a87>

Trade Association and Projects

- Linaro – does ARM upstreaming and projects
 - Still going strong!!
- Robot Operating System (ROS) – handles robotics vertical
- Android Open Source Project (AOSP) – handles mobile phone vertical
 - And there are others (e.g. LineageOS, PinePhone, etc.)
- Yocto Project, Buildroot, OpenWRT – build systems for embedded Linux and routers

Linux Foundation projects

- Linux Foundation
 - Automotive Grade Linux (AGL) – handles automotive vertical
 - Civil Infrastructure Platform (CIP) – handles support longevity
 - Core Embedded Linux Project – is shutting down
 - DroneCode – handles drone vertical
 - ELISA – handled issues with safety certification and standards
 - KernelCI – handles automated testing (for upstream)
 - OpenChain – handles issues with entities in the supply chain
 - OpenSSF – working on security throughout OSS ecosystem
 - Yocto Project – build system for embedded OSS (not just Linux)

Automotive Grade Linux (AGL)

- Latest release = Prickly Pike (16.0)
 - Continued evolution of the Flutter support (by Toyota)
 - Flutter reference apps
 - Flutter workspace automation
 - Chromium Embedded Framework (CEF) available as a build option
 - Instrument cluster container improvements and demo
 - Systemd unit based application activation (tech. preview)
- Based on Yocto 4.0 (kirkstone) release
- Rust meta-layer for Yocto Project
 - Enables the latest Rust version to be used with AGL
 - It was upstreamed to Yocto Project as well
- Check out Walt Miner's talk last week at OSSJ for more info

Civil Infrastructure Platform

- Added 6.1-cip kernel version for super-long-term support
 - All CIP kernels are supported for at least 10 years
- Worked on the IEC 62443 security standard
 - International series of standards for cybersecurity in automation and control systems

DroneCode

- Manages several OSS projects in the drone ecosystem
 - PX4 – open source autopilot software
 - QGroundControl – ground control station software
 - MAVLink – lightweight messaging protocol
 - Pixhawk – open hardware standard for flight management units

ELISA

- ELISA = Enabling Linux in Safety Applications
- New 2023 seminar series about various safety-related topics
- Aerospace Workgroup
 - Working to identify challenges to adopting Linux in aerospace
- Medical Workgroup
 - Published methodology to trace specific workloads to identify kernel subsystem dependencies (for medical applications)
- Published ks-nav tool
 - Perform analysis of Linux kernel according to System-Theoretic Process Analysis (STPA) methods
 - See <https://elisa.tech/blog/2023/06/05/diving-into-the-kernel-introducing-ks-nav-tool-set/>
- Created a reference system with Linux, Xen and Zeph
 - Worked on CI tools

OpenChain


- OpenChain – helps with supply-chain issues
- Standards:
 - Open source license compliance standard ISO/IEC 5230
 - Security assurance program
- Finalized ISO/IEC 18974 standard for security assurance
 - Helps organizations define processes to check OSS for known security vulnerability issues (like CVEs, GitHub alerts or package manager alerts)
 - Identifies security process, roles and responsibilities, etc.
 - Should be published as a formal ISO standard by now (or soon)

Outline

Linux Kernel
Technology Areas
Industry News
Community
Scorecards
Conclusions

Conclusions

- Overall – we're doing very well
 - Embedded Linux is widely deployed and functional (billions of devices)
- Core kernel systems are in place to support embedded
 - But new hardware keeps being made
 - We'll always have things to write and upstream
- More investment is needed by the users of Embedded Linux in the infrastructure and community ecosystem
- What do you think is missing to support embedded development?



Thanks!

